

Abstract

A process for successive refinement to arbitrary precision of a supplied interframe motion estimate for a block of pixels, making use of texture mapping. A bounding box which constrains the limits of the motion estimate must be externally provided, typically one square pixel in size, centered on the original motion estimate. The invention recursively subdivides the supplied bounding box into subregions using a quadtree-like subdivision. A pixel-wise metric comparing the difference between the original block and a prediction for the motion estimate corresponding to each subregion is used to select a particular subregion from the subdivision for further refinement. The prediction is obtained by texture mapping from the target image using the motion estimate corresponding to the center of the subregion. The precision of the refined motion estimate is controlled by bounding the number of refinement steps. Each refinement step provides a doubling of precision in each of the horizontal and vertical directions.